INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 13 Control in Chemical Plants

CHAPTER 1

Introduction: Recycle, Equilibrium and Energy in Chemical Plants

Pages 1–23

Edited by John D. Hedengren

How to cite this chapter. "Introduction: Recycle, Equilibrium and Energy in Chemical Plants," in *Control in Chemical Plants*, John D. Hedengren, Ed., vol. 13 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 1, pp. 1–23.

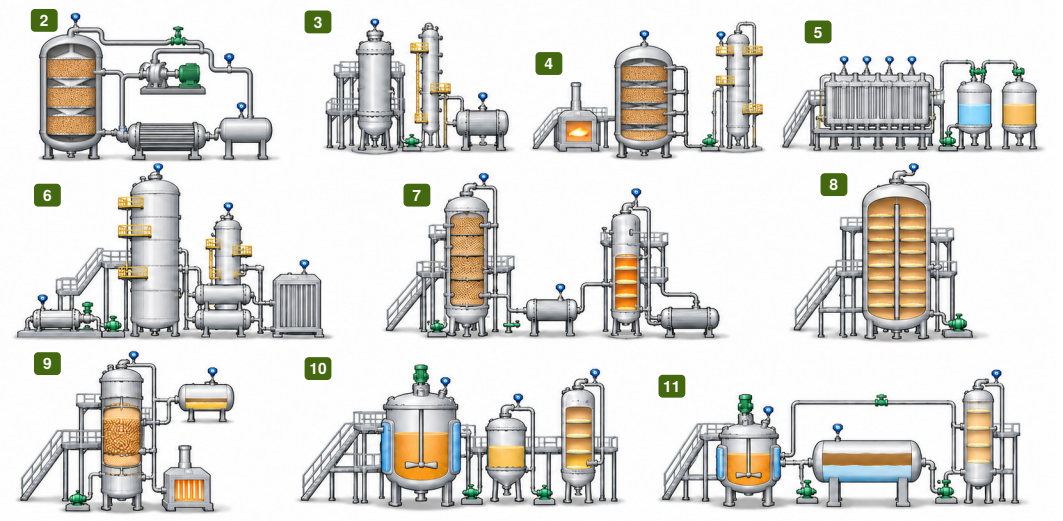
The page numbers in this file are those of the printed volume.

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v13-chemical-plants/

Copyright © 2026 Advanced Process Solutions, LLC. All rights reserved. Produced by Advanced Process Solutions, LLC.

CHAPTER 1

Introduction: Recycle, Equilibrium and Energy in Chemical Plants



The chapters 2 ammonia loop 3 urea loop 4 contact acid plant 5 membrane electrolyzer 6 air separation 7 amine capture 8 dividing-wall column 9 reactive column 10 base oil plant 11 reactor and optimizer

Notice that the top row holds the reactors and cells of the heavy chemicals, the middle row three separations, and the bottom row a reactive column and two reactors, each followed by its separation. Illustration, not to scale.

IN PLAIN TERMS **Plants that send most of their material around again**

The main idea. A chemical plant turns air, natural gas, salt water and sulfur into goods such as fertilizer, acid and chlorine. In a synthesis loop, most of what enters the reactor has been there before, because the part that did not react is sent back around.

The challenge. Many reactions stop partway, at a balance that heat and pressure can shift. Sending material around again makes the plant slow and sensitive. What a stream holds is often known only from a late laboratory test.

Why it matters. Good control decides how much product each unit of energy buys and keeps the plant inside its limits.

THREE WAYS TO READ THIS CHAPTER

Explore	7 min	This page, Section 1.2 with its timeline and two History boxes, Figure 1.2, the Did you know? boxes in Sections 1.4.2, 1.4.3 and 1.4.5, and Key points in Section 1.8. No equations.
Study	45 min	The whole chapter in order, working each number of the example S-1, then Exercises 1 to 10.
Practice	5 min	This page, Section 1.4.8 with Tables 1.1 and 1.2, Section 1.5 with Figures 1.7 and 1.10, and Section 1.8.

1.1 The Question This Volume Answers

A chemical plant turns a few raw materials into a product by reaction and separation, and it rarely does so in one pass. Unconverted reactants return through a recycle, a solvent circulates between an absorber and a stripper, and the heat of one column boils another. The reactions stop at an equilibrium that temperature and pressure move, and the energy that drives the separations is bought at a price

that changes by the hour. This volume answers one question: how do recycle, equilibrium and energy cost shape the control structure of a chemical plant, and which techniques recur from one plant to the next?

Seven control problems define the industry, and most chapters meet more than one of them.

- *Recycle sets the gains and the time constants.* A loop that converts a fifth of its feed per pass, or a solvent that returns after an hour, answers a handle with a gain and a speed that no single vessel has.
- *Equilibrium caps the conversion.* A reversible exothermic reaction runs faster when hot and stops sooner, so temperature acts on conversion in both directions.
- *Some compositions are inventories.* A reactant fed in slight excess to a loop it cannot leave accumulates, and its ratio integrates every feed error.
- *Energy is shared or priced.* One exchanger may be the reboiler of one column and the condenser of another, a spare handle may move only the energy, and the power may cost more in one hour than the next.
- *Composition is inferred.* Loop ratios, solvent loadings and product purities have no fast transmitter, and they are carried between laboratory results by densities, temperatures and models.
- *Feeds follow the throughput.* Stoichiometry or Faraday's law fixes how much of each feed a given rate needs, so the feeds are set in ratio to the throughput and trimmed.
- *The steady state is not unique, and the optimum is known only through a model.* A converter can go out, a column can sit on more than one branch, and an optimizer moves the plant along slopes its model may have wrong.

Ten chapters answer the question on ten plants, each with an app that the reader runs in a browser. Chapters 2 and 3 are the two synthesis loops of a fertilizer site, ammonia and urea. Chapters 4 and 5 are a contact sulfuric acid plant and a membrane electrolyzer. Chapters 6 and 7 separate air and capture carbon dioxide from a flue gas. Chapters 8 and 9 are a dividing-wall column and a reactive column. Chapter 10 makes a synthetic base oil, and Chapter 11 closes the volume with an optimizer above a reactor whose model is wrong.

A reader who has worked through the volume can do the following on a plant that is new to the reader.

- Find where an adiabatic bed stops on the conversion and temperature diagram, and why its inlet temperature has an optimum.
- Predict the gain and the time constant that a recycle or a purge gives a composition, and recognize a composition that integrates.
- Say what a lean loading loop does to the gain of an absorber, and what lowering the loading costs in energy.
- Count the handles of a column, find the one that moves only the energy, and estimate the energy a thermally coupled arrangement saves.
- Set feeds in ratio to the throughput with a trim, and correct an inferred composition from a late laboratory result.
- Explain why a refitted model can stop an optimizer short of the plant's best point, and how a gradient correction removes the gap.

This chapter develops each of these on one small example and names the chapter that uses it. PID forms and tuning are taken as known from Volume 2, ratio, cascade and feedforward from Volume 3, selectors, overrides and mid-ranging from Volume 4, dead time and estimation from Volume 5, the relative gain array and plantwide inventory control from Volume 6, multiplicity and nonlinear analysis from Volume 7 and predictive control and optimization from Volume 9. The refinery and the petrochemical complex that supply several of these plants are Volume 11 and Volume 12.

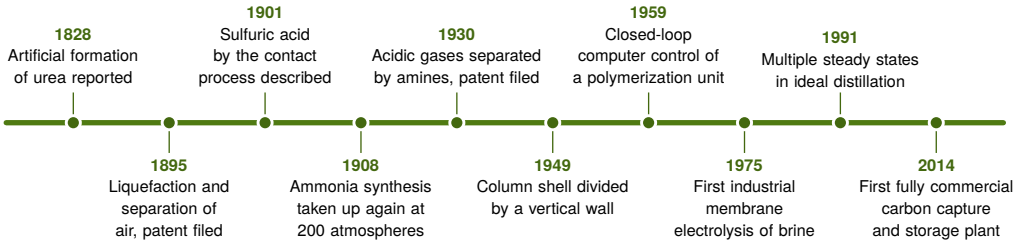


Figure 1.1 Milestones of the chemical plants of this volume and their control. The dates are from [1] (1828), [2] (1895), [3] (1901), [4] (1908), [5] (1930), [6] (1949), [7] (1959), [8] (1975), [9] (1991) and [10] (2014). A patent is dated by its filing where the card says so. Notice that most processes of the volume are older than the computers that now run them. Events are evenly spaced, not to a time scale.

1.2 A Short History

The milestones of Figure 1.1 follow the plants of this volume, from the first heavy chemicals to the separations, the columns and the computers that came to run them.

> ⧖ HISTORY **1905 to 1932: The First Synthesis Loops**
>
> In 1905 Fritz Haber judged ammonia from its elements not worth pursuing, because from the start of red heat no catalyst gave more than a trace of it at normal pressure. He took the problem up again in 1908 at 200 atmospheres, the limit of the compressor he had. His process made and removed the ammonia in a cycle at a constant high pressure, and it led the heat of the exhaust gas to the fresh gas, so that the process supplied its own heat. By 1920 industry had kept the pressure, the circulation and the heat exchange [4]. Urea followed. A mixture of carbon dioxide and ammonia heated in a closed vessel stops at about 40 percent of the theoretical yield, so the excess gases were drawn off at high pressure by heated compressors and condensed again [11]. Chapter 2 and Chapter 3 run both loops.

Air separation began with a patent filed in 1895. Air let down through a valve from a high pressure comes out colder, and the cold air was led back through a long countercurrent coil, about one hundred meters of pipe, to cool the compressed air arriving until it liquefied [2]. A second patent, filed in 1903, rectified the liquid in a column. Compressed gas condensed in a coil in the liquid chamber at the column's foot and boiled that liquid, until pure nitrogen left the top and pure liquid oxygen left the chamber [12]. A coil that condenses one stream and boils another is the shared exchanger of Chapter 6. In 1930 a patent was filed for separating acidic gases with a small group of amines that unite with them when cool and give them up when hot, the solvent circulating in a closed cycle with heat interchangers [5]. That is the solvent loop of Section 1.4.4 and of Chapter 7.

> ⧖ HISTORY **1851 to 1986: Sulfuric Acid and Chlorine**
>
> In 1901 a paper of 47 pages described sulfuric acid and its manufacture by the contact process [3], and Haber later likened the difficulties of scaling up ammonia to those overcome in it [4]. Of the catalysts tested for oxidizing sulfur dioxide only vanadium compounds, platinum and iron oxide proved satisfactory, and vanadium pentoxide on silica, a liquid melt at operating conditions, is now used almost exclusively [13]. Chlorine took longer. An English patent of 1851 made chlorine and caustic soda by electrolysis of brine, and from 1890 the electrodes were separated by a porous diaphragm or replaced by a mercury cathode, unchanged for decades. Stable membranes came at the end of the 1960s, the first industrial membrane plant ran in 1975, and in 1986 Japan converted all its mercury cells [8]. Chapter 4 and Chapter 5 model the plants that came of both.

The columns came next. A column shell divided by a vertical wall through a set of trays was proposed in 1949 for a three-product split, and it is usually named after a theoretical study of 1965 [6]. Reaction and separation in one column have been developed since the early 1920s [14], and in 1991 multiple steady states were shown in a distillation column with ideal thermodynamics [9]. Chapter 8 and Chapter 9 build on both. Computers arrived in 1959, when one achieved closed-loop control of an industrial polymerization unit [7]. The first process computers were supervisory. They adjusted the setpoints of the controllers below them, so the plant kept working when the computer was off [15]. The optimizer of Chapter 11 sits in the same place. The newest milestone is about energy and carbon. In October 2014 the world's first fully commercial carbon capture and storage plant opened at a coal-fired power station [10].

1.3 Notation and Conventions

Temperatures are in °C, differences of temperature in kelvin, and absolute temperatures in a rate or an equilibrium constant in kelvin. Molar flows in the running example are in kmol/h and pressures in bar, written p. A conversion is X, an equilibrium conversion X^*, a gas mole fraction y and a solvent loading x. In the synthesis loop F_M is the makeup, F_P the purge, F the gas fed to the converter and N the gas held in the loop. A recycle is described by a forward gain g_F and the gain g_R of the path that returns. The absorption factor is $\mathscr{A}$ and the stripping factor $\mathscr{S}$. The framework's h is its integration step and its a the controller action, so neither is used for anything else here.

The chapters keep their own symbols, and several collide. λ is a latent heat in Chapter 6 and Chapter 8, an adiabatic rise in Chapter 2, a bias update gain in Chapter 3 and Chapter 10, and a gradient modifier in Chapter 11. The loading of Chapter 7 is α, which Chapter 8 and Chapter 9 use for relative volatility and Chapter 10 for a chain growth probability. r is the hydrogen to nitrogen ratio in Chapter 2, while Chapter 3 writes its ammonia to carbon dioxide ratio as N/C. Chapter 8 splits its liquid by β and its vapor by ρ, and K is an equilibrium constant in some chapters and a filter gain in Chapter 11.

1.4 Recycle, Equilibrium and Energy on One Site

1.4.1 The Plants of This Volume

Figure 1.2 places each chapter's plant among the others. Natural gas is reformed and shifted into a synthesis gas whose hydrogen and nitrogen an ammonia loop converts [16]. Nitrogen can also come from an air separation unit. Ammonia and carbon dioxide make urea, and an amine plant can recover the carbon dioxide from the reformer's flue gas. Sulfur burned in air becomes sulfuric acid, and brine electrolyzed with power becomes chlorine, caustic soda and hydrogen. Among the organic plants, a reactive column makes an ether from a refinery C4 cut and methanol, a dividing-wall column splits aromatics into three products, and a reactor and column make a base oil from decene. The last block is a reactor with an optimizer above it, the benchmark on which Chapter 11 tests real-time optimization.

The chapters take the plants one at a time, and the techniques they share are developed below on one running example, a small synthesis site called S-1, drawn in Figure 1.3. A solvent unit removes an impurity I from a raw gas, and the cleaned gas is the makeup of a loop that makes C from A and B in a two-bed converter. A separator condenses C, a purge removes some of the loop gas, and the rest returns. A column splits the crude C from a light and a heavy byproduct. Every number of S-1 is declared here or computed from these declarations by the script kept with this chapter's record. The numbers are chosen for clear arithmetic and are not taken from any one chapter.

1.4.2 Equilibrium-Limited Reactors

S-1's reaction A + B $\rightleftharpoons$ C runs in the gas at $p = 100\,\text{bar}$ and releases heat. Its equilibrium constant falls as the temperature rises, $K(T) = \exp[B_K(1/T - 1/T_h)]$ in bar^{-1}, with $B_K = 10\,000\,\text{K}$ and $K = 1\,\text{bar}^{-1}$ at $T_h = 525\,\text{K}$. For a feed of equal A and B with an inert fraction y_I, the equilibrium at a conversion

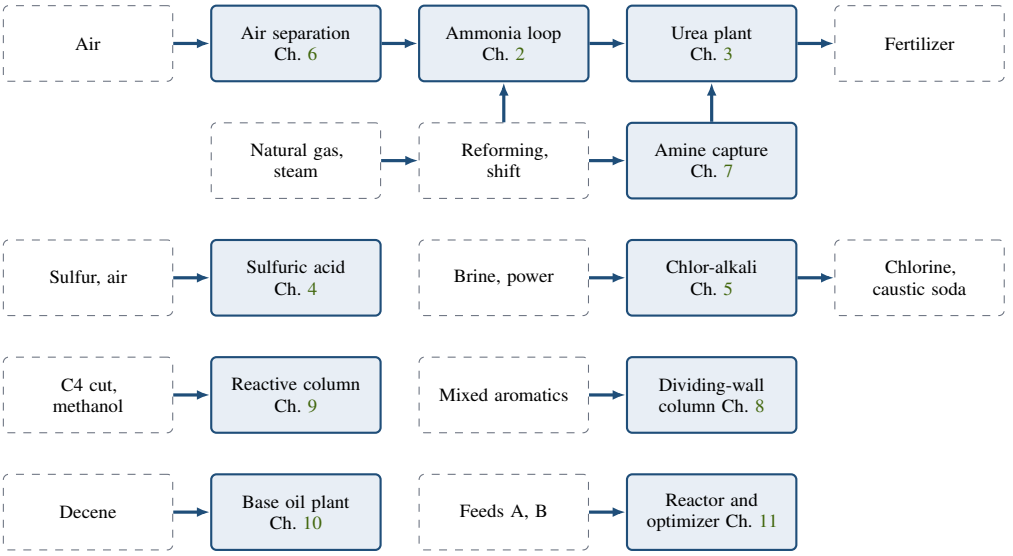


Figure 1.2 The Plants of the Volume. Filled Blocks Are the Plants of Chapters 2 to 11, Dashed Blocks Lie Outside the Volume. The Arrows Show Typical Connections: Nitrogen From Air Separation and Synthesis Gas From the Reformer Feed the Ammonia Loop, and Ammonia With Carbon Dioxide Captured From the Reformer's Flue Gas Feeds the Urea Plant. A Site Need Not Have Every Unit.

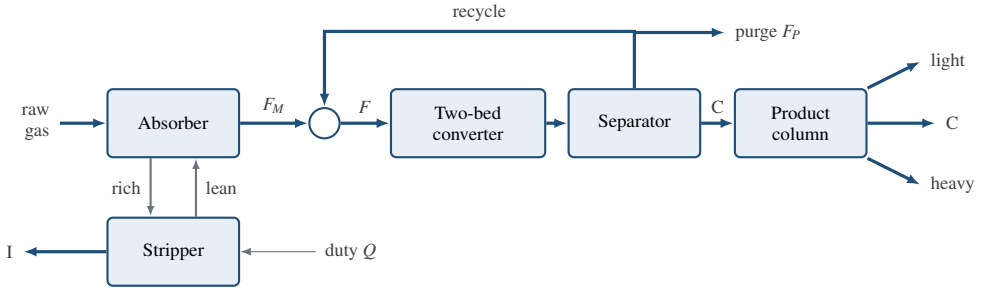


Figure 1.3 The Running Example S-1. A Four-Stage Absorber Removes the Impurity I, Which Is Inert in the Loop, With a Solvent That a Four-Stage Stripper Regenerates. The Loop Converts A + B to C in Two Adiabatic Beds With an Interbed Cooler, Condenses C in the Separator, Purges F_P and Recycles the Rest.

X of A is Equation (1.1).

$$K(T)\,p = \frac{y_C}{y_A y_B} = \frac{X(1-aX)}{a(1-X)^2}, \qquad a = \frac{1-y_I}{2} \tag{1.1}$$

It is a quadratic in X, whose root below one is the equilibrium conversion of Equation (1.2).

$$X^* = \frac{1+2\kappa - \sqrt{(1+2\kappa)^2 - 4\kappa(a+\kappa)}}{2(a+\kappa)}, \qquad \kappa = K(T)\,p\,a \tag{1.2}$$

At 400 °C with 10 % inert, $K = 0.0151\,\mathrm{bar}^{-1}$, $\kappa = 0.680$ and $X^* = 0.345$. A higher pressure raises κ, and more inert lowers it. Chapter 2 solves the same kind of quadratic in closed form for the ammonia equilibrium.

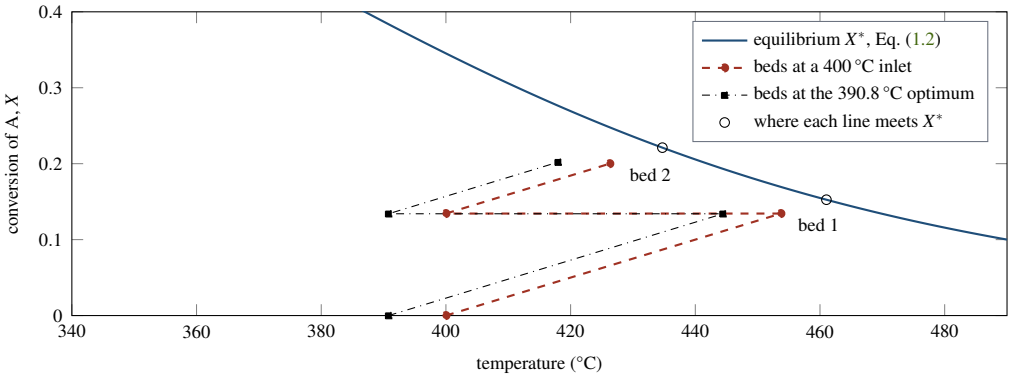


Figure 1.4 S-1's Converter on the Conversion and Temperature Diagram. The Equilibrium Curve (Solid) Falls With Temperature. Each Bed Climbs an Adiabatic Line Toward It and Stops Short (Dashed, Dots), the Cooler Returns the Gas to the Inlet Temperature, and Open Circles Mark Where Each Line Meets the Curve. At the Optimum Inlet (Dash-Dotted, Squares) Both Beds Start Cooler and End Higher.

An adiabatic bed cannot hold its temperature. Each unit of conversion heats the gas by the adiabatic rise ΔT_{ad}, 400 K in S-1, so a bed moves along the straight line of Equation (1.3) and stops where the line meets the equilibrium curve, or short of it.

$$T = T_{in} + \Delta T_{ad}\,(X - X_{in}), \qquad X_{out} = X_{in} + \eta\,(X^*_{line} - X_{in}), \qquad \eta = 1 - \mathrm{e}^{-Da(T_m)} \tag{1.3}$$

where X^*_{line} is the intersection, η the fraction of the way the catalyst gets there, and Da a Damköhler number that grows with the mean bed temperature T_m with an activation temperature of 15 000 K, 1.9 at 700 K. Figure 1.4 draws S-1's converter on this diagram [17]. Both beds enter at 400 °C. Bed 1 would stop on the curve at $X^* = 0.153$ and 461 °C, gets 88 % of the way and leaves at $X = 0.134$ and 454 °C. The interbed cooler returns the gas to 400 °C, and bed 2 climbs a parallel line to $X = 0.200$. One bed alone would give 0.134, and two beds without the cooler 0.152, because the second bed would start next to the curve.

The inlet temperature has an optimum because temperature acts both ways. A hotter inlet speeds the catalyst and lowers the curve it is heading for. S-1 converts 0.193 per pass at 410 °C, 0.170 at 430 °C, and at most 0.2020 at 390.8 °C, 1.0 % more than at the design inlet. Below the optimum the catalyst is too slow, above it the equilibrium too low. Pressure moves the whole picture: at 90 bar the converter gives 0.189 and at 110 bar 0.210. So does the inert. At 12 % inert the per-pass conversion falls to 0.1975 and at 20 % to 0.187, which is how the purge of the next section reaches the reactor.

Chapter 2 runs three quench beds on the Haber equilibrium. Its first bed reaches 82.8 % of the way to equilibrium, and its converter has an optimum near 480 °C of hot spot, worth 11 % of production, against an operating 505 °C. Chapter 4 draws four beds live on this diagram, at 3.20 K per point of conversion, and its first bed stops on the curve at 65.77 %. Removing the product between beds starts the next bed on a new curve. S-1 would reach 0.25 instead of 0.20 if C were taken out after bed 1, and the interpass absorber of Chapter 4 raises its conversion from 97.910 to 99.837 % the same way. The urea reactor of Chapter 3 approaches an equilibrium conversion that rises with its N/C ratio and peaks in temperature, and Chapter 9 removes its ether by distillation as it forms.

> ✦ DID YOU KNOW? **Most Sulfur Becomes Acid**
>
> The world produced an estimated 85 million metric tons of sulfur in 2024, most of it as a result of processing fossil fuels. In the United States about 90 percent of the sulfur consumed that year was in the form of sulfuric acid [18].

A bed that heats its own feed through an exchanger adds a feedback the diagram does not show. The ignited state can then be lost, as in Chapter 2, where a quench of 15.3 % keeps the converter lit and 15.4 % puts it out [19]. Chapter 4 finds an ignition floor near a 395 °C inlet. The loop gain of such a heat exchanger is derived in Volume 11.

1.4.3 Synthesis Loops and Recycle

S-1's loop takes $F_M = 1000\,\mathrm{kmol/h}$ of makeup carrying 1 % of I and purges $F_P = 100\,\mathrm{kmol/h}$. Each mole of C uses one mole each of A and B, so with the purge fixed the loop makes $(F_M - F_P)/2 = 450\,\mathrm{kmol/h}$ of C. At a per-pass conversion of 0.20 the converter needs 2250 kmol/h each of A and B. The separator returns 4000 kmol/h of gas, of which 3900 kmol/h is recycled, 3.9 times the makeup. The loop holds $N = 500\,\mathrm{kmol}$ of gas, seven and a half minutes of the separator's flow. The converter of Section 1.4.2 was computed with the loop gas's inert, although the makeup dilutes it to 8.2 % at the converter inlet, a difference the example leaves out.

An inert can leave only through the purge. With the loop well mixed at constant inventory, its balance is Equation (1.4).

$$N\frac{\mathrm{d}y_I}{\mathrm{d}t} = F_M z_I - F_P y_I, \qquad y_I^{\infty} = \frac{F_M}{F_P} z_I, \qquad \tau_I = \frac{N}{F_P} \tag{1.4}$$

where z_I is the inert in the makeup. Both results are ratios of a large number to a small one. S-1's loop holds 10 % inert from a 1 % makeup, and its time constant is 5 h, forty passes. Chapter 2 measures the same structure: a makeup with 1 % inert and a purge of 8.4 % of the makeup hold 11.90 % in the loop with a time constant of 5.1 h. If the makeup carries 1.2 % inert, S-1 either runs at 12 % and loses 1.3 % of its per-pass conversion, or purges 120 kmol/h to stay at 10 % and makes 2.2 % less C.

> ✦ DID YOU KNOW? **A Fertilizer Made From the Air**
>
> The world's ammonia plants produced an estimated 150 million metric tons of nitrogen in 2024. In the United States natural gas was the main feedstock, and about 88 percent of the ammonia made went to fertilizer [18].

The same result holds for any recycle. If a variable answers a handle with a gain g_F in one pass, and a fraction g_R of the change returns to the inlet, the loop settles at the gain and the time constant of Equation (1.5) [20].

$$K_{ss} = \frac{g_F}{1 - g_R}, \qquad \tau_{ss} = \frac{\tau}{1 - g_R} \tag{1.5}$$

where τ is the time constant of one pass, for a first-order forward path and a return path that is a pure gain. For S-1's inert, the separator gas is the pass, $g_F = F_M/4000 = 0.25$ and $g_R = 1 - F_P/4000 = 0.975$, so both the gain and the time constant are multiplied by forty, to 10 and 5 h. Three regimes recur in the volume. A return gain near one makes a variable large and slow, the snowball effect of reactor and separator plants [21]. A return gain of exactly one, with no exit at all, turns the variable into an integrator. A negative return gain, a recycle that undoes part of the change, makes the settled gain smaller than the gain seen before the recycle returns.

A reactant fed in excess is the integrator case. Suppose the front end feeds B by a meter that reads 1 % high, so the makeup brings 495 kmol/h of A and 490.05 kmol/h of B. The reaction uses them equally, and the excess of A in the loop gas, $\Delta = y_A - y_B$, obeys Equation (1.6).

$$N\frac{\mathrm{d}\Delta}{\mathrm{d}t} = (F_{M,A} - F_{M,B}) - F_P\,\Delta, \qquad \Delta^{\infty} = \frac{F_{M,A} - F_{M,B}}{F_P} \tag{1.6}$$

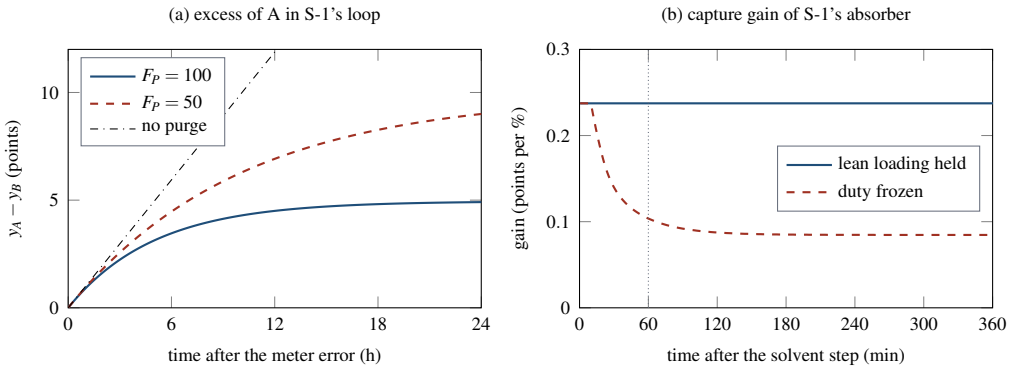


Figure 1.5 Two Recycles of S-1. (a) The Excess of A in the Loop After the B Meter Starts to Read 1 % High, With the Design Purge (Solid), Half the Purge (Dashed) and No Purge (Dash-Dotted). (b) The Apparent Gain of the Absorber's Capture to a 10 % Solvent Step, With the Stripper Duty Frozen (Dashed) and With the Lean Loading Held (Solid). The Dotted Line Marks One Lap of the Solvent.

The excess settles at 4.95 points with the same 5 h time constant as the inert, as shown in Figure 1.5a. The loop gas then holds A and B in the ratio 1.116 against the makeup's 1.010, an error 11.5 times as large. Halving the purge doubles both the error and its time constant. With no purge the excess grows by 0.99 points an hour and never stops. Chapter 2 measures an amplification of 8.0 for a 5 % error in the hydrogen to nitrogen ratio and finds that it is not symmetric, because the real reaction responds to the excess. In Chapter 3 the excess ammonia has no exit, and the N/C ratio integrates at 0.0180 an hour for a fresh ratio 0.02 too high. A ratio held between two measured feeds is not a ratio held in the loop.

The makeup is the only flow that can set the loop's throughput, since every mole of C made and every mole purged entered through it. The design decision is which variable the makeup valve holds [22]. If it holds the loop pressure, the inventory is regulated and a disturbance appears as a change in production. If it holds production, the pressure carries the disturbance, and an override must take the valve back before the pressure reaches its limit. Both arrangements regulate the inventory with a flow that enters it [23]. Chapter 2 offers both. A loss of catalyst activity there costs 17 t/d in one arrangement and raises the pressure to 158.3 bar, 6.7 bar inside the limit of its low selector, in the other.

1.4.4 Absorption and Stripping

S-1's absorber washes I out of 1000 kmol/h of raw gas holding 5 % of it. For a dilute gas with a straight equilibrium line $y = mx$, the outlet of an n-stage counter-current absorber follows the Kremser relation of Equation (1.7) [24].

$$y_{out} = mx_L + (y_{in} - mx_L)\,\varphi(\mathscr{A}), \qquad \varphi(\mathscr{A}) = \frac{\mathscr{A} - 1}{\mathscr{A}^{n+1} - 1}, \qquad \mathscr{A} = \frac{L}{mG} \tag{1.7}$$

where x_L is the lean loading of the solvent entering at the top, L the solvent flow and G the gas flow. The stripper obeys the same relation with the roles of the phases exchanged. It leaves the fraction $\varphi(\mathscr{S})$ of the rich loading in the lean solvent, where the stripping factor $\mathscr{S}$ grows with the vapor the reboiler duty raises. At the design, $\mathscr{A} = \mathscr{S} = 1.25$ with four stages each, so $\varphi = 0.122$, the solvent flow is 1250 kmol/h, the lean loading 0.0044 and the rich loading 0.0365, and the makeup leaves with 1.0 % of I. The absorber removes 80 % of it.

The solvent is a recycle. A larger solvent flow raises $\mathscr{A}$ and captures more at once, but with the duty fixed the same vapor must strip more solvent, so $\mathscr{S}$ falls and the lean loading rises a lap later. Figure 1.5b shows a 10 % solvent step on S-1, whose lap is 60 min. For ten minutes the capture rises

0.237 points per percent of solvent, and with the duty frozen the gain decays to 0.085 as the leaner solvent stops arriving. By Equation (1.5) the return gain is $g_R = 1 - 0.237/0.085 = -1.8$. A lean loading loop that moves the duty to hold x_L cuts the return path, $g_R = 0$, and the gain stays at 0.237. It asks for 7.4 % more vapor, 4.3 % more for each mole of I removed.

Chapter 7 measures this behavior on an amine plant whose lap is 70 min [25]. With the reboiler duty frozen its capture gain falls from 0.177 at ten minutes to 0.040 after eight hours, $g_R = -3.4$, and with the lean loading loop in service it stays at 0.177. A capture loop tuned on the early gain and run with the duty frozen is therefore 4.4 times too weak, and one tuned on the settled gain is 4.4 times too strong once the inner loop is closed. The inner loop is there because it makes one gain out of two.

A leaner solvent costs more than its share of energy. Lowering S-1's lean loading in four equal steps of a fifth of the design value takes 140, 188, 280 and 533 kmol/h more stripping vapor, and the last step buys the same 1.6 points of capture as the first. Chapter 7 pays 38, 43, 51, 65, 90 and 137 MW for six equal steps of loading, because the stripper boils more water to carry less carbon dioxide. The Kremser relation also solves the top of the high-pressure column in Chapter 6, where the heavy key is a few tens of parts per million.

1.4.5 Heat-Integrated and Reactive Columns

A column with a condenser and a reboiler has four flows to move: reflux, distillate, boilup and bottoms. Two hold the drum and sump levels, and two remain for two product purities. A column with a side draw and a liquid split, such as a dividing-wall column, has six handles for three purities and two levels, so one handle is spare. Once the purities are held, the material balances fix the products, and the spare handle can only redistribute the internal flows. It moves the energy and nothing else, and the energy it moves has a minimum.

The minimum is bounded by Underwood's method [26]. For a saturated liquid feed with relative volatilities α_j and mole fractions z_j, the roots θ of Equation (1.8) give the minimum vapor of each sharp split.

$$\sum_j \frac{\alpha_j z_j}{\alpha_j - \theta} = 0, \qquad V_{min} = F \sum_{j \in D} \frac{\alpha_j z_j}{\alpha_j - \theta} \tag{1.8}$$

where the second sum runs over the components in the distillate and θ is the root between the two key volatilities. S-1's crude product is 100 kmol/h of a light, C and a heavy in equal parts with $\alpha = 4$, 2 and 1. The roots are 2.756 and 1.244, the light split needs 107.2 kmol/h of vapor and the heavy split 136.6 kmol/h. Two columns in the direct sequence need $107.2 + 100.0 = 207.2\,\mathrm{kmol/h}$. A thermally coupled arrangement in one shell needs only the larger of the two minima [27], 136.6 kmol/h, a saving of 34.1 %. Chapter 8 finds 30.2 % for its benzene, toluene and o-xylene feed with the same bound. With finite stages its boilup reaches a minimum of 5.83 MW at one liquid split and costs 17.7 % or 21.4 % more at splits a little above and below, and the minimum moves with the feed.

Thermal coupling can also join two columns through one exchanger. In Chapter 6 the nitrogen condensing at the top of the high-pressure column boils the oxygen in the sump of the low-pressure column, across 1.894 K. The driving force is fixed by the duty, so a pressure change in one column must be matched in the other. With the Clausius–Clapeyron slope $\mathrm{d}T/\mathrm{d}p = RT^2/(\lambda_v p)$ on each side, where λ_v is the latent heat, the two pressures move as Equation (1.9) states.

$$\frac{\mathrm{d}p_{HP}}{\mathrm{d}p_{LP}} = \frac{T_{reb}^2\, \lambda_{N_2}\, p_{HP}}{T_{cond}^2\, \lambda_{O_2}\, p_{LP}} \tag{1.9}$$

With the chapter's temperatures and latent heats, 95.36 K and 5000 J/mol at 5.6 bar against 93.47 K and 6700 J/mol at 1.4 bar, the ratio is 2.87. The chapter measures 3.11 bar for each bar of low-pressure column with the bubble points of the real compositions, so pinching a valve on the low-pressure column loads a compressor that feeds the other one.

> **DID YOU KNOW? The Electricity Behind a Tonne of Oxygen**
>
> A typical cryogenic air separation unit uses more than 200 kilowatt-hours of electricity per tonne of oxygen, against a minimum thermodynamic work of around 51 [28]. In Chapter 6 that electricity drives the main air compressor, whose power the controllers cap.

A reactive column removes its product as it forms, which pushes an equilibrium-limited reaction past the limit a reactor would reach. The coupling of reaction and separation can also give one set of handles more than one steady state [9, 29]. Chapter 9 finds three stable states at a duty of 42.5 MW, converting 98.22, 90.05 and 94.89 %. Its stripping-tray temperature peaks at 143.36 °C at 43.15 MW, so 143.13 °C is read at both 43.02 and 43.4 MW. That is an input multiplicity [30]. The process gain from duty to the tray is positive below the peak and negative above it, and a loop whose action was chosen for one side drives the duty the wrong way on the other. The chapter shows such a loop raising the duty without limit on the low-conversion branch.

1.4.6 Ratio Control and Inferred Compositions

When stoichiometry fixes the feeds, the throughput sets each of them. A ratio station multiplies a measured rate by a ratio and adds the output of a trim that reads the composition, as in Equation (1.10).

$$F_{B,SP} = R\,F_{A,m} + u_{trim} \tag{1.10}$$

The ratio is a feedforward that moves the feed the instant the rate moves, and the trim's integral removes what the ratio got wrong. S-1's meter error of Section 1.4.3 is such an unknown. The trim must add 4.95 kmol/h of B, and it sees its own moves through the loop's gain of 0.01 per kmol/h and its 5 h time constant, so it is a slow loop on a slow plant. The membrane electrolyzer of Chapter 5 shows the same split on a fast plant. Faraday's law fixes every product from the current, and with the feeds in ratio to the current and trimmed, a load change keeps the anolyte within 0.28 g/L of its setpoint. With the ratios alone it settles 2.0 g/L low, and with the trims alone it swings 4.88 g/L. Chapter 4 sets its air and its dilution water in ratio to the metered sulfur, and Chapter 9 sets its methanol in ratio to the isobutene fed.

The compositions that matter most are rarely measured fast. A model carries them between laboratory results from secondary measurements [31, 32], and each result corrects the model's bias as in Equation (1.11).

$$b_k = b_{k-1} + \lambda_b\left(y_{lab,k} - \hat{y}(t_{s,k}) - b_{k-1}\right) \tag{1.11}$$

where $\hat{y}(t_{s,k})$ is the model's value at the instant the sample was drawn and $0 < \lambda_b \leq 1$. During S-1's meter error a sample of loop gas drawn after one hour and reported an hour later would, if compared with the model at the report, put 0.73 points into the bias of a model that is right, 45 % of the excess at that time. Chapter 3 infers its N/C ratio from the reactor's density and temperature and corrects it every 4 h from a laboratory that reports 45 min late, with $\lambda_b = 0.7$. Chapter 8 turns tray temperatures into purities with analyzer trims, and Chapter 10 corrects a fouling viscometer from a laboratory every 75 min. Estimation in general is the subject of Volume 5.

1.4.7 Steady-State Optimization Under Model Mismatch

A real-time optimizer waits for a steady state, fits its model to it and moves the setpoints to the model's optimum. This two-step method is Equation (1.12).

$$\theta_k = \arg\min_{\theta} \left\| \bar{y}_k - \hat{y}(u_k, \theta) \right\|^2, \qquad u_{k+1} = \arg\max_{u} \hat{\phi}(u, \theta_k) \tag{1.12}$$

where u are the setpoints, $\bar{y}_k$ the measured steady state, θ the fitted parameters and $\hat{\phi}$ the model's profit. If the model's structure is right, the method converges to the plant's optimum. If it is not, the method stops where the fitted model's slope is zero, which need not be where the plant's is [33].

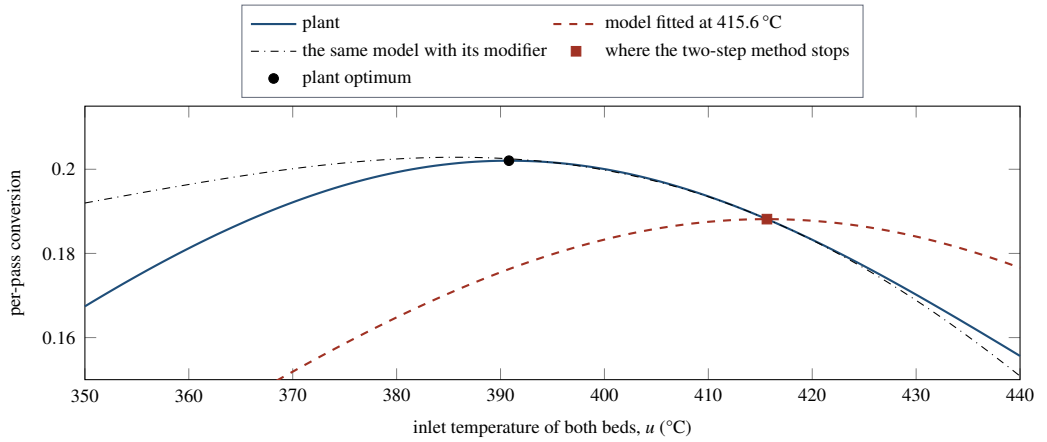


Figure 1.6 Two-Step Optimization on S-1. The Plant's Per-Pass Conversion (Solid) and the Model Fitted Where the Two-Step Method Stops (Dashed) Agree in Value There, but the Model Is Flat and the Plant Falls. Adding the Measured Slope Difference (Dash-Dotted) Moves the Model's Maximum Toward the Plant's Optimum (Dot).

S-1's optimizer sets the inlet temperature of both beds to maximize the per-pass conversion, which at a fixed circulation sets how much makeup the loop can take. Its model is the converter of Section 1.4.2 with an equilibrium constant 43 % too high at every temperature, and it fits only the catalyst activity. As shown in Figure 1.6, from the design inlet of 400 °C, where the plant makes 99.0 % of its best, the method walks to 415.6 °C and stays there, at 93.1 %. The fitted model matches the measured conversion there and predicts a slope of zero, while the plant's slope is −0.0011 per kelvin. A model that reproduces the measurements is not good enough to optimize with, because the optimizer moves along the slopes.

Modifier adaptation keeps the model and corrects it at each run with the measured error and the measured slope [34], as in Equation (1.13).

$$\phi_m(u) = \hat{\phi}(u) + \varepsilon_k + \lambda_k^{\mathsf{T}}(u - u_k), \qquad \lambda_k = \nabla\phi_p(u_k) - \nabla\hat{\phi}(u_k) \tag{1.13}$$

where ε_k is the difference in value and λ_k the difference in gradient, often filtered. The corrected model has the plant's value and slope at u_k, so a point where its slope vanishes is a point where the plant's does. From 415.6 °C with exact plant slopes, S-1's optimizer reaches 99.7 % of the best at its first run and 99.99 % by its fourth, close to the optimum at 390.8 °C. The cost is the plant gradient, which a real plant gives only through deliberate moves and noisy differences.

Chapter 11 measures exactly this on a reactor whose model lacks one of the plant's three reactions. Its two-step optimizer stops at 92.83 % of the best profit while its model predicts more than the plant earns, and modifier adaptation with exact gradients reaches 100.000 % from its tenth run. The other chapters optimize in smaller ways. Chapter 5 plans the current on the hourly price of power with a linear program, Chapter 7 schedules the regeneration energy on the same kind of price, and Chapter 8 moves its spare handle toward the valley's minimum. Chapter 10 backs its reactor off a viscosity index limit that the laboratory sees late. Optimization in general is the subject of Volume 9.

1.4.8 Recycle, Equilibrium and Energy in Each Plant

The three themes of this chapter appear in every plant, though not every plant has all three. Table 1.1 collects what Equation (1.5) predicts and the chapters measure for the plants with a recycle. Table 1.2 lists, for each chapter, what recycles, what equilibrium limits, where the energy goes and the control structure the chapter implements [35]. Chapter 12 adds the measured results and asks how recycle, equilibrium and energy cost shaped each structure.

Table 1.1 What a Recycle Does to a Gain and a Time Constant, in S-1 and in the Chapters

Plant and variable	What returns, how it leaves	Effect
S-1, inert	loop gas, purge 2.5 % of the separator gas	$g_R = 0.975$: 10 % in the loop from 1 %, $\tau_l = 5$ h, forty passes
S-1, excess of A	loop gas, purge	1 % feed error, 11.5 times in the loop ratio, 5 h
S-1, capture	solvent, lean loading	$g_R = -1.8$: 0.237 points per percent at once, 0.085 settled with the duty frozen
2, inert	loop gas, purge 8.4 % of makeup	11.90 % from 1 %, 5.1 h
2, H:N ratio	loop gas, purge	5 % low in the makeup, 8.0 times in the loop, 63 % in 3.2 h
3, N/C ratio	excess ammonia, no exit	$g_R = 1$: integrates 0.0180 an hour for a fresh ratio 0.02 high
7, capture	solvent, lean loading	$g_R = -3.4$: 0.177 at ten minutes, 0.040 after eight hours, 0.177 with the inner loop

Table 1.2 For Each Chapter, What Recycles, What Equilibrium Limits, Where the Energy Goes, and the Structure Implemented

Ch.	Recycle	Equilibrium	Energy	Structure
2	Loop gas, circulation 5.95 times the makeup	Haber, three quench beds	Converter heats its own feed	Makeup on pressure or on production with a low selector, hot spot on quench, slow inert and ratio loops or MPC
3	Unconverted ammonia and carbon dioxide	Conversion rises with N/C	Condenser raises low-pressure steam	Ratio on the flows or on an inferred N/C with laboratory bias, MPC
4	None in the gas	Four beds, double absorption	Bypasses set the bed inlets	Four inlet loops, air and water in ratio to sulfur with trims, MPC on the inlets
5	None	None, Faraday's law	Power on an hourly price	Feeds in ratio to the current with trims, pressure difference cascade with feedforward, price schedule
6	None, one shared exchanger	Vapor and liquid	Compressor near its motor rating	Pressure on the guide vanes with a power override, averaging level, purity loops, MPC
7	Solvent, 70 min lap	Loading against partial pressure	Reboiler duty on a price	Capture on circulation over a lean loading loop, economic supervisor, MPC
8	None	Vapor and liquid	Spare split moves only the duty	Three tray-temperature loops with analyzer bias, split loop, feed feedforward, MPC
9	None	Etherification on catalytic trays	Duty decides the branch	Tray-temperature loop on duty, methanol in ratio, branch monitor, supervisor
10	None	None	Yield against viscosity index, not energy	Viscosity on the cut, mid-ranging to the reactor, bias update, one-way index cap, MPC
11	None	None, irreversible	Jacket, not priced	Three regulatory loops, two-step or modifier adaptation, economic MPC

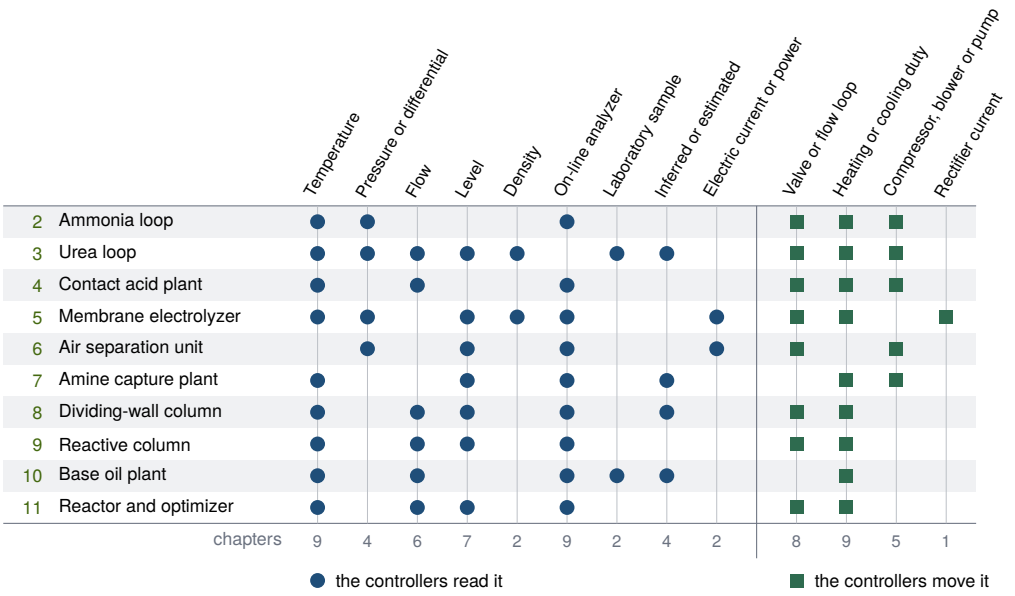


Figure 1.7 Sensors and final elements by chapter. A round mark means that the chapter's controllers or operator read the device, a square mark that they move it. A mark is set only where the chapter's variables table, its instrument text or its section on technology in practice names the device or its reading, so a blank says that the model leaves the device out and not that a real plant lacks it. Notice that only two chapters wait for a laboratory and only one moves an electric current.

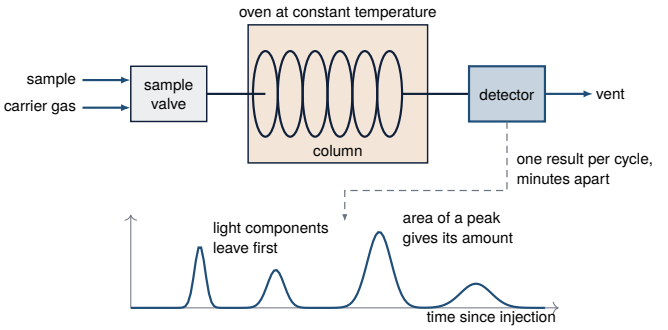
1.5 Instruments and Final Elements in This Volume

Every loop in this volume begins at a sensor and ends at a final element, the device that changes a flow of material or of energy. Figure 1.7 lists what the controllers of each chapter read and what they move. A temperature and an online analyzer are each read in nine of the ten chapters, while a laboratory sample is read in two and a composition or property inferred from other measurements is held in four. A heating or cooling duty is moved in nine chapters and a valve or flow loop in eight. Five chapters move a compressor, a blower or a pump, and one moves the current of a rectifier.

The analyzers set the pace of the composition loops. A chromatograph separates the species of a sample in time, as drawn in Figure 1.8, so cycles of minutes are not uncommon and some dead time is inevitable [36]. The chromatograph of Chapter 2 reports the inerts and the ratio every 15 min, and the analyzers of Chapter 8 sample every 10 min and report a cycle later. Chapter 3 has no analyzer for its ratio at all. It infers the ratio from a density and a temperature and corrects it from a laboratory that reports 45 min after each sample, drawn every 4 h.

Most final elements are valves, and many of them sit on steam, cooling water or a jacket. A positioner compares the stem position with the controller's signal and applies as much air to the actuator as it needs, a small cascade of its own, while a valve with excess friction is unresponsive to small reversals of its signal, most of all without a positioner [36]. Figure 1.9 draws both parts. Large machines carry the energy of several plants. The pressure loop of Chapter 6 acts on the guide vanes of the main air compressor under a power cap, and the rectifier of Chapter 5 sets the current from which Faraday's law fixes every product.

Figure 1.10 places the layers of the volume on their platforms. Regulatory loops act about every second, predictive control every minute and optimization every hour to every day [37, 38]. A distributed control system hosts the loops, and its supervisory level sends setpoints to the field controllers [39]. The apps run their loops every 5 s in Chapter 5 up to every 0.5 min in Chapter 7, their



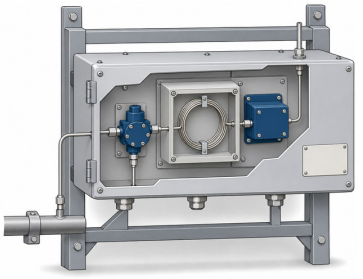
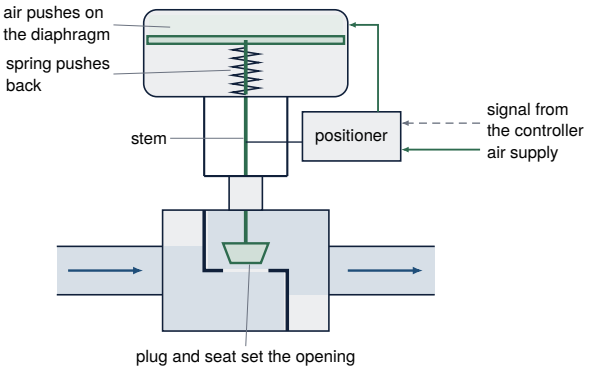


Figure 1.8 Principle of a process gas chromatograph, with a view of the device as installed. Notice that one sample gives one result at the end of the cycle, so a loop on this instrument acts on a value that is always old.



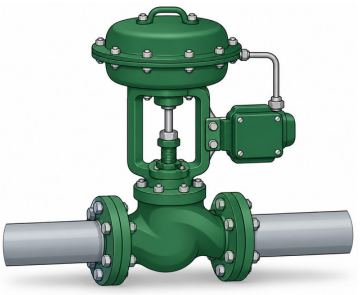


Figure 1.9 Principle of a globe control valve with its spring-and-diaphragm actuator and positioner, with a view of the device as installed. Notice that the positioner closes a loop on the stem position, so the controller's output sets a position, not an air pressure.

predictive controllers every 1 min in Chapter 6 up to every 15 min in Chapter 2, and the optimizer of Chapter 11 no sooner than an hour after its last move.

1.6 The Simulator Framework

Every chapter of this book has an app that runs in a web browser, collected with the others at https://apmonitor.com/apps/. About 240 apps exist across the series. All but a few are built on one shared framework, so the engine, the PID algorithm, the trends and the data export behave the same way in every chapter. The exceptions are standalone studios, such as the cascade studio and the two Temperature Control Lab apps, which follow the same ideas in their own code. This section describes the framework once and states where the standalone apps differ. Table 1.3 lists what a reader meets.

1.6.1 The Simulation Engine

Each app carries a plant model written as ordinary differential equations $\frac{\mathrm{d}x}{\mathrm{d}t} = f(x,u,d,p)$ in the states x, the inputs u, the disturbances d and the parameters p. The engine integrates the model with the classical fourth-order Runge–Kutta method on a fixed step h set for the fastest dynamics of the app. Some apps supply their own step, such as explicit Euler for a discontinuous right-hand side or an exact update for a model driven by events.

Three clocks run at separate rates. The plant advances every step h. The controller executes every control period Δt, equal to h unless the app declares a slower controller, and holds its outputs between

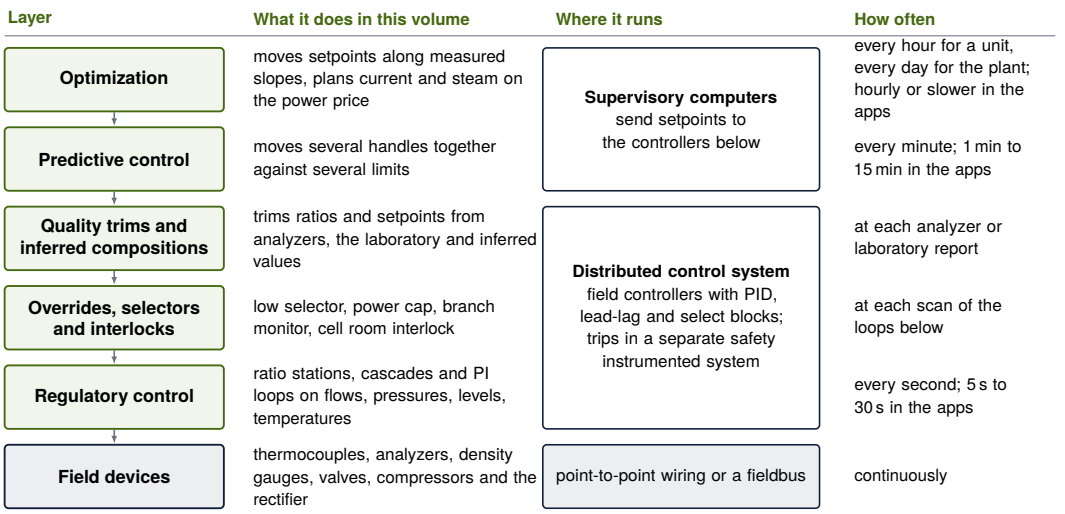


Figure 1.10 The control layers of this volume, the platform that hosts each and how often each acts. The periods of the regulatory, predictive and optimization layers are those of [37, 38], the platforms those of [37, 39], and the ranges in the apps those of the chapters. Notice that each layer sends setpoints to the layer below and acts less often than it.

Table 1.3 What Every App Provides. Standalone Apps Are Described in Their Own Chapters.

Element	What it does
Run, Pause, Reset	Start, hold and restart the run. Reset keeps the reader's settings.
Speed	Simulated time per second of real time. It changes the pace and never the result.
Manual and Auto	Hand each actuator to the reader or to the controller, without a bump.
Scenario list	Numbered cases, each with a briefing, a completion rule and a reload button.
Control panel	Setpoints, manual outputs, tuning, options and disturbances, at a Beginner or an Advanced level.
Trends	SP and PV above the MV for each loop, with limits, event marks and a cursor.
Scorecard	IAE, overshoot, settling time, MV travel and time at a limit for each loop.
Learn card	Predict, Try It, two hints and Explain, opened by the reader in that order.
Show Model	Equations, live variable values, editable parameters and assumptions.
Export CSV	Every logged sample of the run as a text file with a header row.
Instructions	A separate page with background, an operating guide and the model's limitations.

executions. The log takes one sample at a fixed interval of simulated time for the trends, the scorecard and the exported file. The order of work in one step is shown in Figure 1.11. The controller also executes once when a scenario is loaded, so every run starts from an initialized controller.

Drawing the screen is kept apart from these clocks. A speed of 5 means five units of the app's time base per second of real time. Speed, frame rate and pauses change how fast a run is shown and never what it computes, so a phone and a workstation produce the same trajectory and the same file.

Noise comes from a deterministic pseudorandom generator. Its seed is declared by the app, or by a scenario that needs its own noise record, and it is restored whenever a scenario is loaded or Reset is pressed. The framework gives the reader no seed setting, so *default seed* in this book means the seed the app ships with. A run repeats bit for bit when the same scenario is loaded and the same settings are made before Run. A change made by hand during a run lands on a different step each time, which is one reason numeric answers carry a tolerance. The standalone simulators draw noise from the browser's own generator, so a noisy run there repeats in its statistics and not sample for sample.

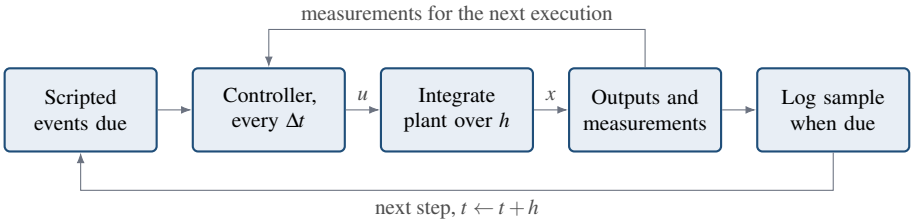


Figure 1.11 Work of the Engine in One Integration Step

Run and Pause start and hold the clock. Reset restarts the run at $t = 0$, paused, and keeps the reader's settings: every setpoint, tuning value, option, model parameter, disturbance slider and mode the reader changed is put back before the plant and the controller are initialized. Scripted events replay, and a one-shot disturbance button pressed during the run is cleared. The reload button beside the scenario list restores the scenario's own settings instead.

1.6.2 The Velocity-Form PID Algorithm

One PID implementation serves the PID loops of every framework app. It is the velocity (incremental) form: each execution adds the change in the proportional and derivative terms to a running sum, and the integral is a separate term, as shown in Equations (1.14) to (1.17) for execution k and control period Δt.

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c\,e_k, \qquad D_k = -a\,K_c\,\tau_D\,\frac{\mathrm{PV}^f_k - \mathrm{PV}^f_{k-1}}{\Delta t} \tag{1.14}$$

$$S_k = S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \tag{1.15}$$

$$v_k = u_b + S_k + I_{k-1}, \qquad u_k = \mathrm{sat}(v_k) \tag{1.16}$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I}\,e_k + \frac{u_k - v_k}{T_t}\right]\Delta t \tag{1.17}$$

where $a = \pm 1$ is the controller action, PV^f is the measurement after an optional first-order filter, u_b is the output bias, S is the running sum of proportional and derivative increments, v is the output before limiting, u is the output sent to the actuator, I is the integral term and T_t is the tracking time constant, equal to τ_I unless a loop sets another value. The function sat clamps v_k to the actuator range and then applies a rate limit if the loop declares one.

The limit acts on the emitted output u_k only. The sum S_k is never clamped, so it equals $P_k + D_k$ and the output before limiting is the position form $u_b + K_c e_k + D_k + I_{k-1}$. Anti-windup touches the integral alone. While the output is limited, the last term of Equation (1.17) drives I toward the value that reproduces the output applied, which is back-calculation [40]. With no integral action, Equation (1.17) is skipped. Clamping the running sum would instead discard each increment that crosses a limit and keep each one that returns, so symmetric noise would push the output off the limit even in a proportional-only loop.

The derivative acts on the measurement, so a setpoint step gives a proportional kick and no derivative kick. A change of K_c acts on the whole error at the next execution, as in the position form. The action $a = +1$ suits a process whose PV rises with the MV, a reverse-acting controller in the ISA convention, and $a = -1$ gives the direct-acting one.

In Manual the controller keeps executing and sets $I_k = u_{\mathrm{man}} - u_b - S_k$, so the output it would send equals the manual value u_{man}. The first automatic output differs from it by one period's change in P and D. In Auto the engine copies each controller output to the manual station, so a switch to Manual is also bumpless. The standalone studios with PID loops use velocity-form code of their own, with

the limit on the output and anti-windup on the integral alone, in which a gain change applies only to later increments.

1.6.3 The Predictive Control and Estimator Options

More than 100 apps offer model predictive control (MPC) as an option beside the structure they open with, which stays the default. Volume 9 develops the theory. In most of them the controller plans with a small linear model in discrete time as shown in Equation (1.18).

$$x_{k+1} = A\,x_k + B\,u_k + E\,w_k + c, \qquad y_k = C\,x_k \tag{1.18}$$

where w is a measured or scheduled disturbance and c is a constant left by linearization. The model comes from a step test at the design point, from the app's balances linearized there or again at every decision, or from the model the reader fits. It never reads the true states or parameters of the run, so a parameter the reader changes in Show Model is model mismatch, as on a real plant.

At each decision the controller minimizes the objective of Equation (1.19) over a horizon of N steps.

$$J = \sum_{k=1}^{N} (y_k - r_k)^{\top} Q\,(y_k - r_k) + \sum_{j} \Delta u_j^{\top} R\,\Delta u_j + \sum_{i} \rho_i\, s_i^2 \tag{1.19}$$

where r is the setpoint, Q weights the tracking error, R is the move suppression on each input change Δu_j and ρ_i is the price of a slack $s_i \geq 0$ that lets an output band be broken when it cannot be held. Actuator limits are hard constraints. The quadratic program is solved in the browser by a dual active-set method [41]. The first move is applied, and the plan is made again at the next decision.

The plan starts from an estimate of the model state. About half of these apps use a steady-state Kalman filter on the model with an added disturbance state, which absorbs what the model did not predict and removes steady offset without an integral term [42]. Others correct the model output with a filtered bias to the same end. About a dozen apps also offer moving horizon estimation (MHE), which solves the estimation problem over a window of recent samples and can hold bounds that a Kalman filter cannot, such as a concentration that is never negative [43].

With the MPC selected, the tuning fields of the loops it replaces leave the Control panel. In most of these apps the Advanced level shows up to three settings, each a multiple of the shipped value: the prediction horizon, the move suppression and the band price. The trend draws the accepted plan ahead of the present time, and a list of constraint prices shows which limits bind. A plan that is not optimal, not finite or outside its limits is rejected, and a fallback, in most apps the shipped controller, holds the actuator until a plan succeeds. The Model Fitting and Tuning Studio builds the same kind of controller from the reader's own data in its tab, MPC from the fit.

1.6.4 Scenarios, Progress and Data

An app teaches through scenarios, the numbered cases at the top of the Control panel, nine to thirteen in most apps. Loading one sets the model parameters, the controller settings, the disturbances, the mode, the initial state and any scripted events. The information button beside the list opens a briefing in four parts: what the case is, why it matters, what to do and one detail for the practicing engineer. The Learn card holds the lesson back, with Predict and Try It before the run, hints on request and Explain afterward.

Each scenario has a completion rule that the app judges from the run on the simulated clock: a run length, the scripted events passed, a setting made by the reader, a state held for a stated time or a target met. Choosing a case earns nothing. A completed case is marked in the list, and a badge row shows the count and a level, Explorer after one case, Practitioner at half and Completed at all. Progress is kept in the reader's browser with no account. A certificate of progress or of completion can be issued from it, as a record of self-paced work and not an assessed credential.

One case, usually the last, is the graded challenge. It states a target, such as an error band to hold through a sequence of upsets. Changing the model parameters or the disturbances voids the attempt in most apps, and a record made with Reveal suggested tuning says so.

Show Model opens the equations, live values of the variables, the editable parameters and the assumptions. The trends plot SP and PV above the MV for each loop, with actuator limits and marks at setpoint moves, disturbances, mode changes and scripted events. The scorecard reports for each loop the integral of the absolute error (IAE), overshoot, settling time, MV travel and share of time at a limit, with the app's own measures, and it does not combine them into one score. IAE and travel accumulate from Reset, so two runs are compared at the same clock reading.

Export CSV saves the log as a text file with a header row of column names that carry their units and one row per logged sample. In most apps the time column is in seconds, minutes or hours according to the run length. The log holds the most recent 60 000 samples unless the app sets another size. Apps with the MPC option begin the file with comment lines that start with #, which a script must skip.

An app exercise in this book names the scenario by its number and name as the app shows them, every setting that differs from the scenario's defaults, the run length in simulated time, the seed (the default seed unless stated), what to read and where (a trend, a scorecard tile or a column of the exported file) and a tolerance on a numeric answer. Settings are made before Run unless the exercise gives a time.

1.7 Working With AI Agents While Studying

A reader of this book is likely to have an AI assistant within reach, and the apps make its work easy to test. This section gives the editor's view of how to use one while working through a chapter. In short, the assistant may do the algebra and the scripting, and the reader keeps the checking.

What to delegate. Tasks with a clear input and a result that can be tested are good candidates. An assistant can derive the linearization of a chapter's balance at an operating point and the transfer function that follows from it. It can write a short script that reads the exported CSV file and fits a first-order-plus-dead-time model to a step test. It can propose a first tuning from that fit with a named rule, convert units, explain a term in a model or outline a report. The assistant knows nothing about the app until it is told. Give it the equations and parameter values from Show Model, the column names of the exported file and the scenario that produced the data, as described in Section 1.6.4. Ask it to state its assumptions and the units of every result.

What to check, and how. An assistant returns a fluent answer whether or not the answer is right, so each result needs a test that it could fail. Four tests cover most of the errors a chapter will produce.

- *Units.* Every term of an equation carries the same units, and a process gain is in PV units per MV unit. A gain computed with the valve as a fraction is 100 times the gain of a loop whose valve is in percent.
- *Signs.* Decide from the physics which way the PV moves when the MV rises. Compare that with the sign of the fitted gain and with the controller action of Section 1.6.2.
- *Steady state.* Set the derivatives to zero. The model must return the operating point that Show Model reports, and a linearization must give zero deviation there.
- *A run on the app.* Load the named scenario at the default seed, enter the proposed numbers and run. The trend and the scorecard decide. A tuning that reads well on paper and holds the valve at a limit for half the run has failed.

A fitted model gets one more test: compare it with data it was not fitted to, such as a second step in the opposite direction. When a check fails, report the failure to the assistant with the evidence and ask for the cause, then check the correction the same way.

Table 1.4 Roadmap of the Volume

Ch.	Plant	The lesson	Level	Time
2	Ammonia synthesis loop	Recycle sets the gains, and makeup holds pressure or rate.	advanced	3 h
3	Urea synthesis loop	The loop's ratio integrates every feed error and is inferred.	introductory	2 h
4	Contact sulfuric acid plant	The coolest inlet converts most and sits nearest the floor.	intermediate	2 h
5	Membrane electrolyzer	The current is the load, so every feed follows it, and the gas side is faster than feedback.	introductory	2 h
6	Air separation unit	One exchanger couples two columns, so no valve is local.	intermediate	3 h
7	Amine capture plant	An inner loop makes the gain a constant, not the loop faster.	advanced	3.5 h
8	Dividing-wall column	The spare handle moves only the energy, along a valley.	intermediate	2 h
9	MTBE reactive column	The same handles hold more than one steady state.	advanced	2 h
10	Synthetic base oil plant	Two handles reach one viscosity with two different oils.	introductory	2.5 h
11	Reactor with an optimizer	A fitted model stops short, and the slopes decide.	advanced	2.5 h

How the exercises train the habit. One exercise in each chapter is marked *Review the agent*. It presents a tuning, a model or a claim as an AI assistant might propose it, with a flaw, and asks the reader to find the flaw in about five minutes. The flaws are ordinary ones: a unit, a sign, a steady state that does not close, a result measured on the wrong scenario or at a different run length, or a limit that was ignored. Each is found with one of the tests above. The purpose is not distrust. An engineer who can say why an answer is right can hand an assistant much larger tasks and still stand behind the result. That larger picture, in which agents do most of the routine work and the engineer supervises and checks, is the subject of Chapter 12.

1.8 Chapter Roadmap

Table 1.4 lists the chapters with the lesson each adds to this one. The level says how much of Section 1.4 a chapter assumes. The time is an estimate for reading the chapter once and running each case of its app once, at about eight minutes a case.

Chapters 3, 5 and 10 form a first course: a ratio that integrates, feeds that follow a current, and two handles on one quality with a late laboratory. They rest on Sections 1.4.3 and 1.4.6. Chapters 4, 6 and 8 add an equilibrium-limited reactor train, thermal coupling and a column's energy valley. Chapters 2, 7, 9 and 11 combine a recycle that moves gains and time scales, more than one steady state, or an optimizer with a wrong model, and are best left until last. A reader who follows the fertilizer chain can read Chapters 6, 2, 7 and 3 in that order.

KEY POINTS

- In a synthesis loop most of what enters the reactor has been there before, so the recycle sets how large and how slow each response is.
- A ratio held between two measured feeds is not a ratio held in the loop, because a reactant that cannot leave integrates every feed error.
- A lean loading loop does not make the absorber faster. It makes one gain out of two.
- The spare handle of a column with a side draw moves only the energy, and the energy it moves has a minimum.
- A model that reproduces the measurements is not good enough to optimize with, because the optimizer moves along the slopes.

Where else this applies. Recycle and inventory control return in Volume 6, ratio control in Volume 3, estimation in Volume 5 and optimization in Volume 9.

References

[1] F. Wöhler, “Ueber künstliche bildung des harnstoffs,” *Annalen der Physik*, vol. 88, no. 2, pp. 253–256, 1828.

[2] C. Linde, “Process of producing low temperatures, the liquefaction of gases, and the separation of the constituents of gaseous mixtures,” U.S. Patent 727,650, filed July 9, 1895, granted May 12, 1903, 1903. [Online]. Available: https://patents.google.com/patent/US727650A/en

[3] R. Knietsch, “Ueber die Schwefelsäure und ihre Fabrication nach dem Contactverfahren,” *Berichte der deutschen chemischen Gesellschaft*, vol. 34, no. 3, pp. 4069–4115, 1901.

[4] F. Haber, “The synthesis of ammonia from its elements,” Nobel Lecture, 2 June 1920, 1920, nobelprize.org/uploads/2018/06/haber-lecture.pdf, read November 2026.

[5] R. R. Bottoms, “Process for separating acidic gases,” U.S. Patent 1 783 901, Dec. 2, 1930, filed October 7, 1930.

[6] E. A. Wolff and S. Skogestad, “Operation of integrated three-product (Petlyuk) distillation columns,” *Industrial & Engineering Chemistry Research*, vol. 34, no. 6, pp. 2094–2103, 1995.

[7] Computer History Museum, “Computers in control,” CHM Revolution, online exhibition, section Real-Time Computing, accessed 18 November 2026. [Online]. Available: https://www.computerhistory.org/revolution/real-time-computing/6/130

[8] M. Paidar, V. Fateev, and K. Bouzek, “Membrane electrolysis—history, current status and perspective,” *Electrochimica Acta*, vol. 209, pp. 737–756, 2016.

[9] E. W. Jacobsen and S. Skogestad, “Multiple steady states in ideal two-product distillation,” *AIChE Journal*, vol. 37, no. 4, pp. 499–511, 1991.

[10] M. Bui, C. S. Adjiman, A. Bardow, E. J. Anthony, A. Boston, S. Brown, P. S. Fennell, S. Fuss, A. Galindo, L. A. Hackett, J. P. Hallett, H. J. Herzog, G. Jackson, J. Kemper, S. Krevor, G. C. Maitland, M. Matuszewski, I. S. Metcalfe, C. Petit, G. Puxty, J. Reimer, D. M. Reiner, E. S. Rubin, S. A. Scott, N. Shah, B. Smit, J. P. M. Trusler, P. Webley, J. Wilcox, and N. Mac Dowell, “Carbon capture and storage (CCS): the way forward,” *Energy & Environmental Science*, vol. 11, no. 5, pp. 1062–1176, 2018.

[11] C. Bosch, “The development of the chemical high pressure method during the establishment of the new ammonia industry,” Nobel Lecture, 21 May 1932, 1932, nobelprize.org/uploads/2018/06/bosch-lecture.pdf, read November 2026.

[12] C. Linde, “Apparatus for producing pure nitrogen and pure oxygen,” U.S. Patent 795,525, filed October 22, 1903, granted July 25, 1905, 1905. [Online]. Available: https://patents.google.com/patent/US795525A/en

[13] European Commission, “Reference document on best available techniques for the manufacture of large volume inorganic chemicals: Ammonia, acids and fertilisers,” European Commission, Integrated Pollution Prevention and Control, Tech. Rep., Aug. 2007, eippcb.jrc.ec.europa.eu, read November 2026.

[14] A. A. Kiss and M. Jobson, “Taking reactive distillation to the next level of process intensification,” *Chemical Engineering Transactions*, vol. 69, pp. 553–558, 2018.

[15] M. Pitt, “Taking a look back at control: Part 2,” *The Chemical Engineer*, Oct. 2023, published online 26 October 2023. [Online]. Available: https://www.thechemicalengineer.com/features/taking-a-look-back-at-control-part-2/

[16] M. Appl, “Ammonia, 2. production processes,” in *Ullmann’s Encyclopedia of Industrial Chemistry*. Wiley-VCH, 2011.

[17] M. J. King, W. G. Davenport, and M. S. Moats, “Oxidation of SO_2 to SO_3—equilibrium curves,” in *Sulfuric Acid Manufacture*. Elsevier, 2013, pp. 125–133.

[18] U.S. Geological Survey, *Mineral Commodity Summaries 2025*. Reston, VA: U.S. Geological Survey, 2025, nitrogen (fixed)–ammonia, pp. 128–129, and Sulfur, pp. 172–173.

[19] J. C. Morud and S. Skogestad, “Analysis of instability in an industrial ammonia reactor,” *AIChE Journal*, vol. 44, no. 4, pp. 888–895, 1998.

[20] M. M. Denn and R. Lavie, “Dynamics of plants with recycle,” *The Chemical Engineering Journal*, vol. 24, no. 1, pp. 55–59, 1982.

[21] W. L. Luyben, “Snowball effects in reactor/separator processes with recycle,” *Industrial & Engineering Chemistry Research*, vol. 33, no. 2, pp. 299–305, 1994.

[22] R. M. Price, P. R. Lyman, and C. Georgakis, “Throughput manipulation in plantwide control structures,” *Industrial & Engineering Chemistry Research*, vol. 33, no. 5, pp. 1197–1207, 1994.

[23] E. M. B. Aske and S. Skogestad, “Consistent inventory control,” *Industrial & Engineering Chemistry Research*, vol. 48, no. 24, pp. 10 892–10 902, 2009.

[24] M. Souders and G. G. Brown, "Fundamental design of absorbing and stripping columns for complex vapors," *Industrial & Engineering Chemistry*, vol. 24, no. 5, pp. 519–522, 1932.

[25] G. T. Rochelle, "Amine scrubbing for CO_2 capture," *Science*, vol. 325, no. 5948, pp. 1652–1654, 2009.

[26] A. J. V. Underwood, "Fractional distillation of multicomponent mixtures," *Industrial & Engineering Chemistry*, vol. 41, no. 12, pp. 2844–2847, 1949.

[27] I. J. Halvorsen and S. Skogestad, "Minimum energy consumption in multicomponent distillation. 2. three-product Petlyuk arrangements," *Industrial & Engineering Chemistry Research*, vol. 42, no. 3, pp. 605–615, 2003.

[28] M. Variny, D. Jediná, M. Rimár, J. Kizek, and M. Kšiňanová, "Cutting oxygen production-related greenhouse gas emissions by improved compression heat management in a cryogenic air separation unit," *International Journal of Environmental Research and Public Health*, vol. 18, no. 19, p. 10370, 2021.

[29] S. Hauan, T. Hertzberg, and K. M. Lien, "Why methyl tert-butyl ether production by reactive distillation may yield multiple solutions," *Industrial & Engineering Chemistry Research*, vol. 34, no. 3, pp. 987–991, 1995.

[30] L. B. Koppel, "Input multiplicities in nonlinear, multivariable control systems," *AIChE Journal*, vol. 28, no. 6, pp. 935–945, 1982.

[31] B. Joseph and C. B. Brosilow, "Inferential control of processes: Part I. steady state analysis and design," *AIChE Journal*, vol. 24, no. 3, pp. 485–492, 1978.

[32] P. Kadlec, B. Gabrys, and S. Strandt, "Data-driven soft sensors in the process industry," *Computers & Chemical Engineering*, vol. 33, no. 4, pp. 795–814, 2009.

[33] J. F. Forbes, T. E. Marlin, and J. F. MacGregor, "Model adequacy requirements for optimizing plant operations," *Computers & Chemical Engineering*, vol. 18, no. 6, pp. 497–510, 1994.

[34] A. Marchetti, B. Chachuat, and D. Bonvin, "Modifier-adaptation methodology for real-time optimization," *Industrial & Engineering Chemistry Research*, vol. 48, no. 13, pp. 6022–6033, 2009.

[35] J. J. Downs and S. Skogestad, "An industrial and academic perspective on plantwide control," *Annual Reviews in Control*, vol. 35, no. 1, pp. 99–110, 2011.

[36] T. R. Kuphaldt, *Lessons in Industrial Instrumentation*, version 2.33 ed. Open textbook, Creative Commons Attribution 4.0 license, ibiblio.org/kuphaldt/socratic/sinst, 2024.

[37] S. Qin and T. A. Badgwell, "A survey of industrial model predictive control technology," *Control Engineering Practice*, vol. 11, no. 7, pp. 733–764, 2003.

[38] S. Skogestad, "Plantwide control: The search for the self-optimizing control structure," *Journal of Process Control*, vol. 10, no. 5, pp. 487–507, 2000.

[39] K. Stouffer, M. Pease, C. Tang, T. Zimmerman, V. Pillitteri, S. Lightman, A. Hahn, S. Saravia, A. Sherule, and M. Thompson, "Guide to operational technology (OT) security," National Institute of Standards and Technology, Gaithersburg, MD, Special Publication 800-82r3, 2023.

[40] K. J. Åström and L. Rundqwist, "Integrator windup and how to avoid it," in *Proc. 1989 American Control Conference*, 1989, pp. 1693–1698.

[41] D. Goldfarb and A. Idnani, "A numerically stable dual method for solving strictly convex quadratic programs," *Mathematical Programming*, vol. 27, no. 1, pp. 1–33, 1983.

[42] G. Pannocchia and J. B. Rawlings, "Disturbance models for offset-free model-predictive control," *AIChE Journal*, vol. 49, no. 2, pp. 426–437, 2003.

[43] C. V. Rao, J. B. Rawlings, and D. Q. Mayne, "Constrained state estimation for nonlinear discrete-time systems: Stability and moving horizon approximations," *IEEE Transactions on Automatic Control*, vol. 48, no. 2, pp. 246–258, 2003.

1.9 Exercises

1.1 Theory and calculation · 5 min

A gas-phase reaction A + B $\rightleftharpoons$ C is fed with equal A and B, and at the bed temperature $K(T)\,p = 2$, as in Section 1.4.2.

(a) Find the equilibrium conversion of A with no inert in the feed. *Checkpoint:* $a = 0.5$, $\kappa = 1$, $X^* = 0.423$.

(b) Find it again with 10 % inert at the same $K(T)\,p$. *Checkpoint:* $a = 0.45$, $\kappa = 0.9$, $X^{*} = 0.398$.

1.2 Concept check · 2 min

A two-bed adiabatic converter for a reversible exothermic reaction runs both inlets at 400 °C. Raising both inlets to 420 °C lowers the per-pass conversion. Why?

(a) the catalyst slows down when it is hotter, so each bed gets less far along its line

(b) the adiabatic rise per unit of conversion grows with temperature, so each bed heats faster

(c) each bed starts closer to the falling equilibrium curve, so the line meets it at a lower conversion

(d) the interbed cooler can no longer return the gas to the inlet temperature

1.3 Theory and calculation · 4 min

A synthesis loop holds 600 kmol of gas. Its makeup is 1200 kmol/h with 0.8 % of an inert that leaves only through a purge of 80 kmol/h.

(a) Find the settled inert fraction in the loop and its time constant. *Checkpoint:* 12.0 % and 7.5 h.

(b) How long after a step in the makeup's inert does the loop reach 95 % of its new level? *Checkpoint:* About 22.5 h.

1.4 Theory and calculation · 5 min

In S-1 the meter on B starts to read 2 % high, so the makeup brings 495 kmol/h of A and 485.1 kmol/h of B. The purge stays at 100 kmol/h, the makeup's inert at 10 kmol/h and the loop at 500 kmol.

(a) Find the settled excess $y_A - y_B$ and the time constant. *Checkpoint:* 0.099 and 5 h.

(b) With the loop's C production $(F_M - F_P)/2$, find y_A, y_B and the loop ratio y_A/y_B, and compare its error with the makeup's. *Checkpoint:* 0.4995, 0.4005, 1.247 against 1.020, about 12 times.

1.5 Concept check · 2 min

A urea loop sets its ammonia in ratio to the measured carbon dioxide, and the excess ammonia returns through the recycle with no purge. The ammonia meter starts to read 1 % high. What does the ammonia to carbon dioxide ratio in the reactor do?

(a) it falls and keeps falling until something outside the ratio station corrects it

(b) it falls by about 1 % and settles there, as the feeds do

(c) it falls by about ten times 1 % and settles with a time constant of hours

(d) it does not change, because the ratio station holds the measured ratio

1.6 Theory and calculation · 3 min

A circulation step raises an absorber's capture with a gain of 0.20 points per unit before the returning solvent changes, and the solvent loop returns a fraction $g_R = -3$ of the change once it has gone around.

(a) Find the settled gain. *Checkpoint:* 0.050 points per unit.

(b) A capture loop is tuned for the settled gain, and then an inner lean loading loop is closed. By what factor is the capture loop now too aggressive? *Checkpoint:* Four.

1.7 Theory and calculation · 4 min

A three-stage absorber treats a dilute gas with 4 % of a solute. The equilibrium line is $y = mx$, the absorption factor is $\mathscr{A} = 1.5$, and the lean solvent's back pressure is $mx_L = 0.002$. Find the outlet mole fraction and the fraction of the solute removed.

1.8 Theory and calculation · 5 min

S-1's crude product is 100 kmol/h of saturated liquid with a light, C and a heavy in equal parts, at relative volatilities 4, 2 and 1. The Underwood roots are $\theta_1 = 2.756$ and $\theta_2 = 1.244$.

(a) Find the minimum vapor of the light split, with only the light in the distillate, and of the heavy split, with the light and C in the distillate. *Checkpoint:* 107.2 kmol/h and 136.6 kmol/h.

(b) The second column of the direct sequence needs 100.0 kmol/h. Compare the direct sequence with a thermally coupled column. *Checkpoint:* 207.2 kmol/h against 136.6 kmol/h, 34.1 % less.

1.9 Concept check · 3 min

A real-time optimizer refits its model at every steady state, and after a few runs it stops moving. The fitted model reproduces the measured outputs at the final point. What can be concluded?

(a) the plant is at its optimum, because a model that fits the measurements has the plant's slopes
(b) the plant is at its optimum only if the steady-state detector was strict enough
(c) the model's structure must be wrong, because a refitted model never stops
(d) the fitted model's slope is zero there, which says nothing certain about the plant's slope

1.10 Theory and calculation · 4 min

A plant's profit is $\phi_p(u) = -(u-5)^2$, and its model is $\hat{\phi}(u) = -(u-3)^2 + c$, whose only fitted parameter is the offset c.

(a) Where does the two-step method stop? *Checkpoint:* $u = 3$.

(b) From $u_k = 3$, find the gradient modifier with exact plant gradients and the next setpoint of modifier adaptation. *Checkpoint:* $\lambda_k = 4$, $u_{k+1} = 5$.